SACLANTCEN REPORT
serial no: SR- 330

SACLANT UNDERSEA RESEARCH CENTRE REPORT

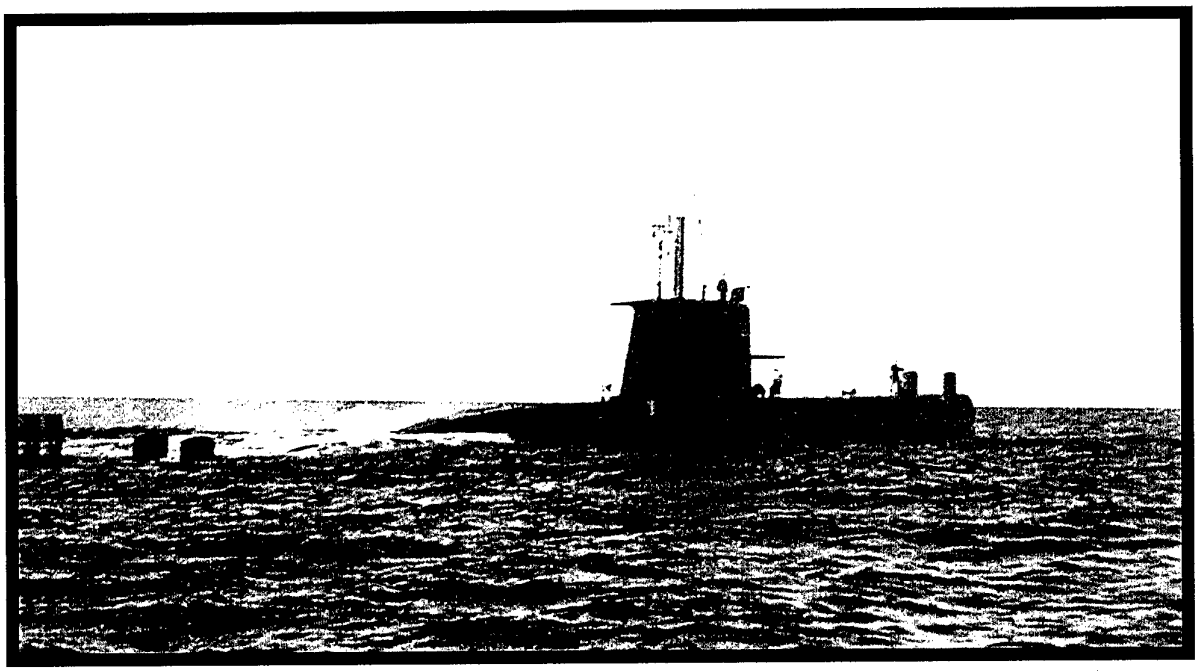

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

NORTH ATLANTIC TREATY ORGANIZATION ORGANISATION DU TRAITE DE L'ATLANTIQUE NORD

SACLANTCEN SR-330

Markov random field segmentation for low frequency active sonar: further results

R. Laterveer

The content of this document pertains to work performed under Project 04-A of the SACLANTCEN Programme of Work. The document has been approved for release by The Director, SACLANTCEN.

Jan L. Spoelstra
Director

intentionally blank page

Markov random field segmentation for low frequency active sonar: further results

R. Laterveer

Executive Summary: Low frequency active sonar has been highlighted by a number of NATO nations as an important component of the next generation of undersea defence systems. The use of low frequencies in a shallow water environment, however, is known to result in a high false alarm rate due to the large number of clutter like returns which can overload automatic tracking and classification algorithms.

The SACLANT Undersea Research Centre is currently investigating techniques to aid in the reduction of these false alarms without a reduction in detection probability.

A previous study described an automatic method of image segmentation based on Markov random field (MRF) modelling to reduce clutter. The method examines detections over both range and bearing and removes small objects which do not exhibit the right signature over beams. Separate detections corresponding to one large object are combined to form a single object. Objects too large to be a submarine can then be removed.

In this report the algorithm is evaluated on additional data. Using Markov random field segmentation to detect large geophysical features and a small submarine simultaneously does not improve performance. Markov random field segmentation could still be used to detect large geophysical objects which can be removed if their geological position remains unchanged or their size is too large to be a submarine. This will be examined in future work.

intentionally blank page

Markov random field segmentation for low frequency active sonar: further results

R. Laterveer

Abstract:

The use of low frequency active sonar in shallow water leads to large numbers of clutter detections.

This high false alarm rate can overload automatic tracking and classification algorithms. Traditional detection algorithms operate on each beam output individually searching for targets at all ranges. However, the target echo and bottom features may extend over several beams, either because a reflector is extended over space or because of the sidelobe structure of the beamformer. This suggests the association of detections over bearing, e.g. apply image processing to the range-bearing sonar data.

A previous study described an automatic method of image segmentation based on a Markov random field (MRF) model to reduce clutter. The segmentation is treated as a labelling problem, assigning to each range-bearing cell either a target or background label, removing small objects which do not exhibit the correct signature over beams. Separate detections corresponding to one large reflector are combined and removed if they are too large to be a submarine .

In this report, the algorithm is evaluated on additional data. On a single ping basis the MRF segmentation shows slightly lower detection performance at the same number of false alarm objects as the Page test detector. MRF segmentation still shows potential for fixed feature removal over multiple pings and will be studied in the feature.

Keywords: clutter reduction ○ detection ○ classification ○ segmentation ○ low frequency active sonar ○ Markov random fields ○ image processing

Contents

1 Introduction 1

2 Markov random fields for active sonar 2

3 Data processing 5

4 Performance 7
4.1 Page test performance 7
4.2 MRF Performance 7

5 Conclusions 11

References 12

1

Introduction

Low frequency active sonar is an important component of the next generation of undersea defence systems. The use of low frequencies in a shallow water environment, however, is known to result in a high false alarm rate due to the large number of clutter like returns. The SACLANT Undersea Research Centre is investigating automatic techniques to reduce false alarms without a reduction in detection probability.

Target detection is usually not a problem, generally the signal to noise ratio (SNR) is sufficient. The large number of false detections makes it difficult to find the target. As the next stage of processing is target/non-target classification using algorithms for tracking and fixed-feature removal, the large number of false targets makes the computational load of such an approach prohibitive.

Automatic processing becomes especially important for broadband sonar. Due to the higher resolution of broadband sonar there is more information, making it difficult or even impossible to present it all to a sonar operator.

The separation of regions of extended background reverberation is perceived to be a crucial first step in categorisation. A segmentation algorithm may be considered as a first stage classifier which allows the estimation of size and position of extended objects and consequently facilitate their elimination and/or classification.

In [1, 2, 3] an image segmentation algorithm based on a Markov random field (MRF) model was proposed. In this report the algorithm will be further analysed and applied to additional data.

In Sect. 2 we briefly discuss the MRF algorithm. Section 3 shows the different processing steps and the way in which we measure the performance. The results are presented in Sect. 4. We end with conclusions and recommendations in Sect. 5.

2

Markov random fields for active sonar

An MRF algorithm for segmentation of active sonar pictures was proposed [1] and refined and evaluated on detection performance [2, 3]. In this section we describe the algorithm.

Background (non-target like) pixels are labelled with a 0 and target-like returns with a 1. The probability that a pixel in our image has a particular value $\lambda \in \{0, 1\}$ is given as $P(\Lambda = \lambda)$, the so called *a priori* probability.

We shall incorporate two basic assumptions into the model

i) that a pixel of a given type is surrounded by pixels of the same type (the usual assumption for segmentation algorithms).

ii) that the background can be modelled as a Rayleigh distribution and that the 'targets' are non-Rayleigh. This will allow the introduction of the concept of a conditional probability or threshold which separates the two distributions.

The Hammersly-Clifford theorem [4] allows us to express the global probabilistic model in terms of a local energy model through the Gibbs distribution:

$$P(w) = \frac{1}{Z} e^{-U(w)/T}, \tag{1}$$

where $Z = \sum_w \big(\exp -U(w)/T\big)$ is called the partition function, T is constant called the temperature, and $U(w)$ is the energy function. Maximising the global probability now corresponds to minimising the total energy.

The construction of the *maximum a posteriori* (MAP) model is implemented by using Bayes' theorem. The *maximum a posteriori* probability is given by [4],

$$\begin{aligned} \hat{\lambda}_{\text{opt}} &= \arg\max_{\hat{\lambda} \in \Omega} P(\Lambda = \hat{\lambda} | Z = z) \\ &= \arg\max_{\lambda \in \Omega} P(Z = z | \Lambda = \lambda) P(\Lambda = \lambda) \end{aligned} \tag{2}$$

Where Z represents the observed value of the pixel, $\hat{\lambda}_{\text{opt}}$ is the optimal estimated label for a particular pixel, $P(Z = z | \Lambda = \lambda)$ is the conditional probability and $P(\Lambda = \lambda)$ is the *a priori* probability.

This formulation allows us to write our problem as a well defined optimality criterion. The optimisation is then performed in an iterative manner. For every pixel in the

image the (locally) optimal label is determined. After several iterations the labelling converges to the global optimum.

It is now necessary to introduce physical arguments to allow us to describe the terms on the right hand side of the above equation. In the *a priori* probability we allow two forms of interaction: local, single pixel dependent terms, and a second order contribution which encompasses an area around the pixel of interest. This second order term is introduced as an Ising model [4] which acts as an homogenising force which, physically, can be considered as stating that background pixels tend to be be surrounded by background pixels. The first order term is introduced to reduce the probability of small objects disappearing. The overall *a priori* energy terms can be written then as

$$U(\Lambda = \lambda) = -\alpha \ln(c_s) - \beta \sum \delta(\lambda_j, \lambda), \tag{3}$$

where α and β are user-defined homogenising parameters and δ is the Kronecker delta operator. The summation is taken over an area surrounding the pixel of interest. We have taken a window three pixels wide (equivalent to three beams) and 5 pixels long (approximately 30 m for the datasets considered). The *a priori* energy can be seen to decrease the overall energy as neighbouring pixels line up.

The homogenising constants are increased at every iteration. This allows acceleration of regularisation without removing small objects too quickly, slowly moving from a conditional probability dominant model to an *a priori* dominant model [5]. We choose $\alpha = k\beta$, where k is a constant, and let β increase at every iteration from 0.5 in steps of 0.1 to 1.

If the pixel under investigation is a '0' in the present iteration then $c_s = [p_{00}\delta(\lambda_s, 0) + p_{10}\delta(\lambda_s, 1)]$, where λ_s is the pixel value after the iteration and p_{10} is the probability of changing from a 0 to a 1. A similar expression exists if the present pixel value is 1. N.B. $p_{00} + p_{10} = 1$ and $p_{01} + p_{11} = 1$. p_{00} is a measure of the likelihood of a background pixel staying a background pixel and is taken to be 0.5. p_{11} is a measure of the likelihood of an object pixel staying as an object. Consequently, this should be set high (typically greater than 0.8). Increasing p_{11} will increase the probability of target detection, but also increase the probability of false alarm. We consider two possible models [3], one with p_{11} fixed over all iterations and one in which after each iteration we determine the maximum likelihood estimate $\hat{p}_{11}$ which maximises the joint probability $\prod P(\lambda_i; p_{11})$

$$\hat{p}_{11} = \frac{n_{11}}{n_{11} + n_{01}}, \tag{4}$$

where n_{ij} the number of transitions from label j to label i and the product is taken over the entire sonar picture. This estimate is then used for the following iteration.

The conditional probability is rather more easily defined [1]. The background and object distributions are well separated (Fig. 1), so we consider a 'conditional threshold'

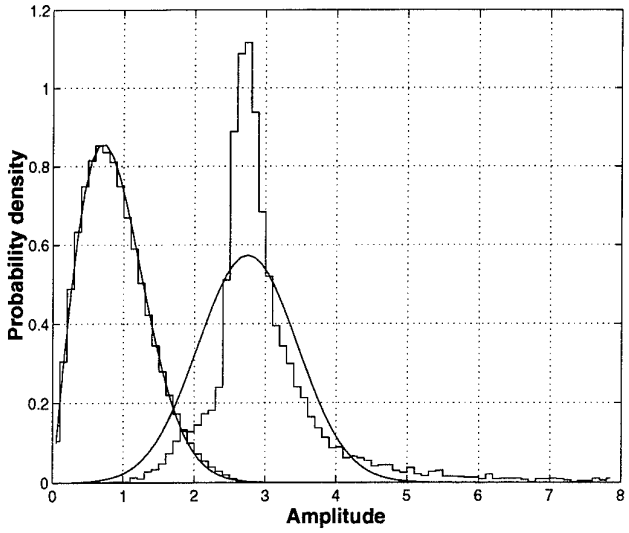

Figure 1 *Histograms of background pixels and Page test detections. The solid lines show a Rayleigh pdf and a noncentral Rayleigh pdf with 8.5 dB SNR.*

z_t and assign probabilities:

$$\text{if} \quad z \le z_t \quad \text{then} \quad P(Z = z|\Lambda = 0) = p \tag{5}$$

$$\text{and if} \quad z > z_t \quad \text{then} \quad P(Z = z|\Lambda = 1) = q. \tag{6}$$

From Fig. 1 we see that p and q can be determined from the cumulative probability densities (CDFs)[3]:

$$p = \frac{\int_0^{z_t} f_0(z)dz}{\int_0^{z_t} \big(f_0(z) + f_1(z)\big)dz} \tag{7}$$

$$q = \frac{\int_{z_t}^{\infty} f_1(z)dz}{\int_{z_t}^{\infty} \big(f_0(z) + f_1(z)\big)dz}, \tag{8}$$

where $f_0(z)$ is the background pdf and $f_1(z)$ is the object pdf.

We model the object pdf with a noncentral Rayleigh distribution. Note that this is just an approximation, the object pixels correspond to different returns with different SNR. Even so the approximation is appropriate. A reasonable fit is a noncentral Rayleigh pdf with 8.5 dB SNR, a good choice for the threshold is $z_t = 2$. With these distributions we get $p = 0.8$ and $q = 0.97$. This is confirmed by the cumulative histograms of the data

3

Data processing

We show results for two data sets, taken south of Sicily. Data set 1 was taken in September 1996 and uses a 2.29 s HFM pulse from 460 Hz to 565 Hz. Data set 2 was taken in June 1995 and uses a 1.2 s HFM pulse from 470 Hz to 590 Hz. Assuming 750 m/s for the two-way sound speed this results in a range resolutions of respectively about 7 and 6 m. Both data sets have around 170 pings with one or more target detection opportunities. The target SNR varies from ping to ping but is, for most pings, strong enough to produce a detection.

The processing steps are illustrated in Fig. 2.

The data were beamformed using 128 hydrophones at 1 m spacing using Hann shading over the array. The beams were equally spaced in cosine-space, and overlapped 3 dB down from their main response axis at 700 Hz. Data set 1 has 86 beams spanning from forward to aft endfire. Data set 2 has 34 beams encompassing 50° around broadside.

Each beam was matched filtered and basebanded so that the centre of the waveform, 512.5 Hz for data set 1 and 530 Hz for data set 2, was shifted to zero Hz.

The beams were normalised using a split window trimmed-mean normaliser [6]. The normalising window consisted of a window of about 1200 m around the sample to be normalised. The samples in the normalising window were ordered and the lower and upper quarter were removed. The remaining samples were used to estimate the power in the auxiliary window. After normalisation the sample frequency is reduced by taking the maximum sample in every range resolution cell resulting in a timeseries with almost independent samples.

Under ideal normalisation the background will have a chi-squared distribution with two degrees of freedom

$$p_{\chi^2_2}(z) = \frac{1}{2} e^{-z/2}.$$

The MRF segmentation is initialised by the Page test detector [7] which is a sequential detection method. For chi-squared distributed data the Page test statistic is a restricted cumulative sum

$$W_n = \max\{0, W_{n-1} + z_n - b\}$$

where b is the bias. The statistic W_n is compared to a threshold h. A detection begins when the threshold is exceeded and ends when it again falls under the threshold and

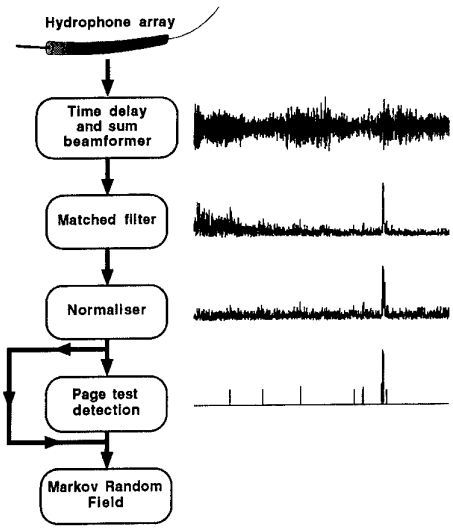

Figure 2 *Processing chain.*

thus also measures the target extend. The threshold h depends on the sampling frequency, i.e. doubling the sampling frequency requires a threshold about 3 dB higher. The bias does not depend on the sampling frequency. The value of the bias can be related to the expected SNR of the target [7].

For each ping a human operator has determined if it contains a target echo with at least 10 dB SNR. The target is then represented by the range/bearing cell which contains the maximum echo peak. This is taken as the ground truth knowledge.

The performance is evaluated after both the Page test detector and the MRF segmentation. The performance measure uses the notion of objects, range/bearing cells which are connected count as one object [8], i.e. all the returns, both over range and beams, due to a single object are treated as one single entity. This procedure automatically takes care of an object's range extend and bearing extend, both physically as well as due to the beampattern's sidelobes.

After processing a target is said to be detected if the maximum echo peak falls within an object and the object's centre of gravity is inside the beam of the echo peak and within 100 m range. The center of gravity is used to prevent very large objects from being counted as valid target detections[1] as they would give no information on the target's position. The performance is shown by plotting the probability of detection against the number of false alarm objects per ping. Note that this is not a ROC curve, as the horizontal axis shows the number of false alarm objects per ping as opposed to false alarm probability per detection bin.

[1]The unlikely case where the target is at the centre of the sonar screen and the object covers the entire sonar screen could still be counted as a valid detection.

4 Performance

4.1 Page test performance

Figures 3 and 4 show the Page test detection results. The different curves are for different values of the Page test bias, the different values of the Page test threshold are represented by the circles. From the figures it is clear that data set 2 possesses stronger target echoes.

For low bias, the detection probability decreases because at too low bias, the Page test detects most of the sonar picture producing large objects which are not counted as valid target detections. When the bias is too big the detection probability drops. The optimal bias seems to be 8 dB, which is used to initialise the MRF segmentation.

4.2 MRF Performance

Figures 5 and 6 show the MRF segmentation results for data set 1. The results for the MRF model with fixed p_{11} are shown in Fig. 5, while the results for the MRF model with adaptive p_{11} are shown in Fig. 6. The dashed line shows the results for the Page test detector with bias 8 dB, a threshold of 6 dB, indicated by the large dot is used to initialise the MRF segmentation. The solid lines show the results for the MRF segmentation, each line is for a different value of $k = \alpha/\beta$, the symbols on the line are for different values of p_{11}. The results of the adaptive p_{11} model are not very dependent on the choice for the initial value of p_{11}, i.e. the results are more robust than the results for fixed p_{11}.

MRF segmentation does reduce the number of false alarm objects by 80% coupled with a reduction in detection probability. Reducing the threshold for the Page test detector, for fixed bias, also reduces the number of false alarm objects, with a smaller reduction in detection probability. The MRF segmentation has about 5% to 7% reduction in detection probability.

Figures 7 and 8 show the MRF segmentation results for data set 2. The results are similar to data set 1. Again the results of the adaptive p_{11} model are more robust than the results for fixed p_{11}. The number of false alarm objects is reduced up to 86% compared to the Page test initialisation. Choosing the wrong segmentation parameters can increase the number of false alarm objects.

For typical number of false alarm objects, the detection probability is reduced by about 7% for data set 1 and 2% for data set 2 (Table 1).

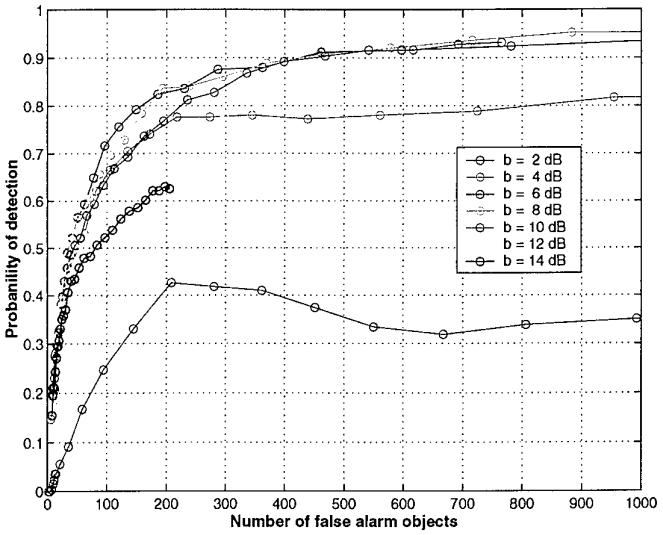

Figure 3 *Page detection performance curve for data set 1. The different curves are for different values of the Page test bias, the different values of the Page test threshold are represented by the circles.*

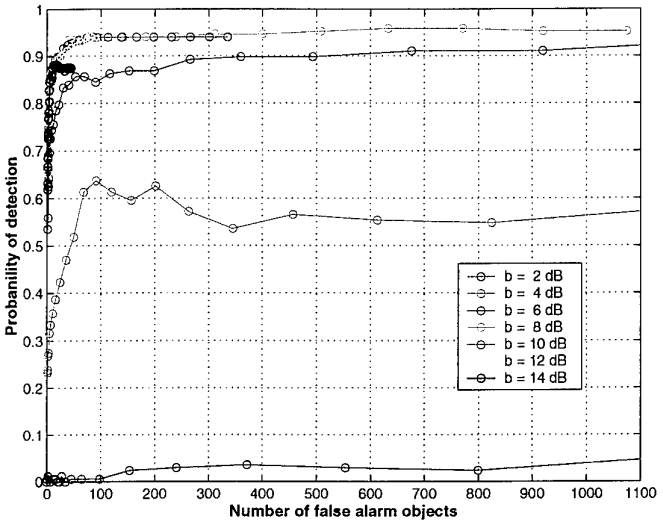

Figure 4 *Page detection performance curve for data set 2. The curves are as in Fig. 3.*

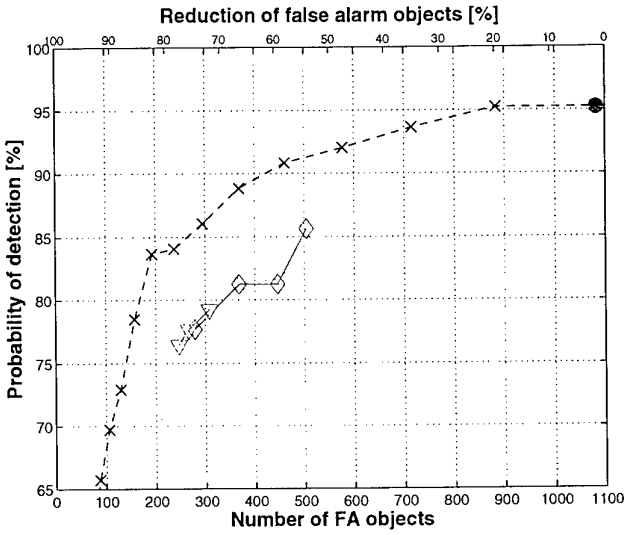

Figure 5 *Performance of MRF model with fixed p_{11} for data set 1. The dashed line shows the results for the Page test detector with bias 8 dB, a threshold of 6 dB, indicated by the big dot, is used to initialise the MRF segmentation. The full lines show the results for the MRF segmentation, each line is for a different value of $k = \alpha/\beta$, the symbols on the line are for different values of p_{11}.*

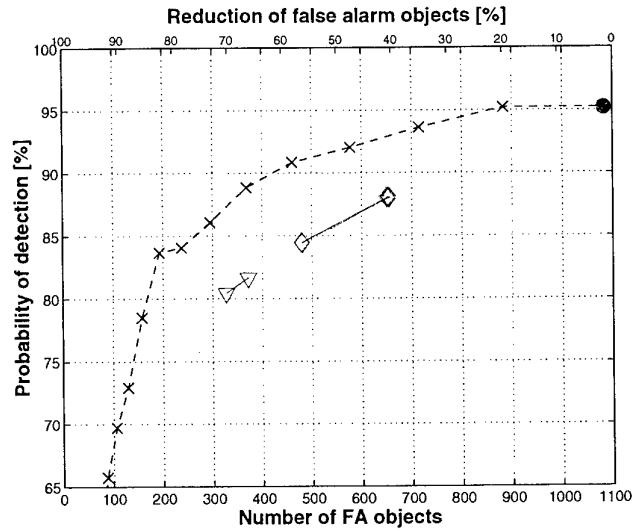

Figure 6 *Performance of MRF model with adaptive p_{11} for data set 1. The lines and symbols are as in Fig. 5.*

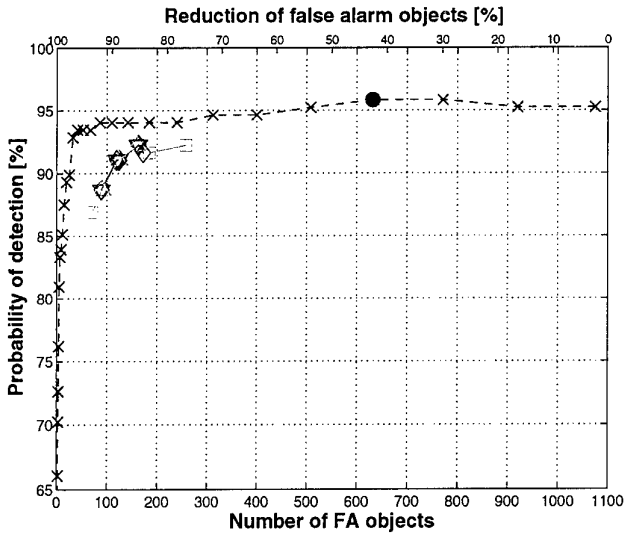

Figure 7 *Performance of MRF model with fixed p_{11} for data set 2. The lines and symbols are as in Fig. 5.*

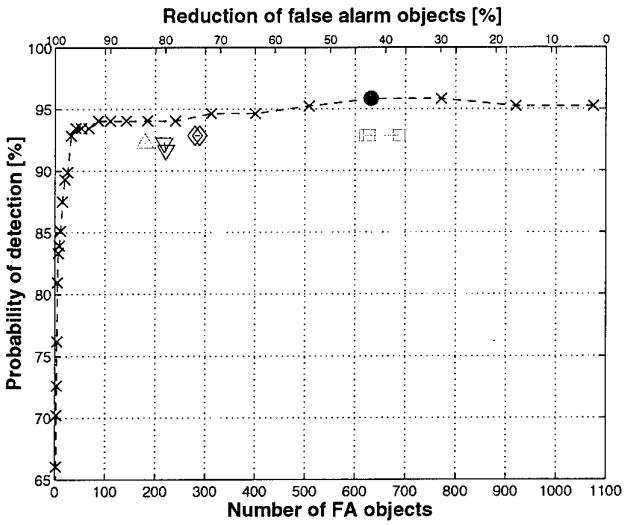

Figure 8 *Performance of MRF model with adaptive p_{11} for data set 2. The lines and symbols are as in Fig. 5.*

	approx. FA	Page test	MRF fixed p_{11}	MRF adaptive p_{11}
data set 1	300	86%	79%	79%
data set 2	200	94%	92%	92%

Table 1 *Comparison of detection probability between Page test detector and MRF segmentation at typical fixed number of false alarm objects.*

5 Conclusions

This report describes a segmentation algorithm based on a Markov random field model using the Page test detector for initialisation. The algorithm was tested on two data sets of about 170 pings each with one or two targets per ping.

For the single ping case, MRF segmentation has no advantage over the Page test detector. MRF segmentation reduces the number of false alarms significantly, but so does reduction of the Page test detector threshold, with a slightly better probability of detection.

The Markov random field segmentation algorithm was designed to merge separate returns from one geophysical object. Detecting large geophysical objects and a small submarine are conflicting demands so it might be difficult to perform both with the same algorithm. Markov random field segmentation still shows potential to detect large geophysical objects which can be eliminated depending on unchanged geological position over ping. In this case submarine detection probability is not relevant.

Traditionally *a priori* information for fixed feature removal is included through a Kalman filter. MRF modelling is an alternative method by extending the range-bearing information with information over ping. This will be studied in future work.

References

[1] Dugelay, S., and Abraham, D. Reduction of low frequency active sonar clutter through image processing. SR-272, SACLANTCEN, 1998.

[2] Laterveer, R., Hughes, D., and Dugelay, S. Markov random fields for target classification in low frequency active sonar. In *Oceans'98*, pages 1274–1278, 1998.

[3] Laterveer, R. Single ping clutter reduction: segmentation using Markov random fields. SR-307, SACLANTCEN, 1999.

[4] Li, S. *Markov Random Field Modeling in Computer Vision.* Springer, 1995.

[5] Besag, J. On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society*, Series B 48:259–302, 1986.

[6] Ghandhi, P., and Kassam, S. Analysis of CFAR processors in nonhomogeneous background. *IEEE Transactions on Aerospace and Electronic Systems*, 24:427–445, 1988.

[7] Abraham, D., and Willett, P. Active signal detection in shallow water using the Page test. SR–252, SACLANTCEN, 1996.

[8] van Velzen, M., and Laterveer, R. Performance measurement in active sonar using object counting. SR-311, to be published, SACLANTCEN, 1999.

Document Data Sheet

Security Classification		Project No.
UNCLASSIFIED		04-A
Document Serial No.	**Date of Issue**	**Total Pages**
SR-330	December 1999	18 pp.

Author(s)

Laterveer, R.

Title

Markov random field segmentation for low frequency active sonar: further results.

Abstract

The use of low frequency active sonar in shallow water leads to large numbers of clutter detections.

This high false alarm rate can overload automatic tracking and classification algorithms. Traditional detection algorithms operate on each beam output individually searching for targets at all ranges. However, the target echo and bottom features may extend over several beams, either because a reflector is extended over space or because of the sidelobe structure of the beamformer. This suggests the association of detections over bearing, e.g. apply image processing to the range-bearing sonar data.

A previous study described an automatic method of image segmentation based on a Markov random field (MRF) model to reduce clutter. The segmentation is treated as a labelling problem, assigning to each range-bearing cell either a target or background label, removing small objects which do not exhibit the correct signature over beams. Separate detections corresponding to one large reflector are combined and removed if they are too large to be a submarine.

In this report, the algorithm is evaluated on additional data. On a single ping basis the MRF segmentation shows slightly lower detection performance at the same number of false alarm objects as the Page test detector. MRF segmentation still shows potential for fixed feature removal over multiple pings and will be studied in the feature.

Keywords

Clutter reduction – detection – classification – segmentation – low frequency active sonar – Markov random fields – image processing

Issuing Organization

North Atlantic Treaty Organization
SACLANT Undersea Research Centre
Viale San Bartolomeo 400, 19138 La Spezia, Italy

[From N. America: SACLANTCEN (New York) APO AE 09613]

Tel: +39 0187 527 361
Fax:+39 0187 527 700

E-mail: library@saclantc.nato.int

The SACLANT Undersea Research Centre provides the Supreme Allied Commander Atlantic (SACLANT) with scientific and technical assistance under the terms of its NATO charter, which entered into force on 1 February 1963. Without prejudice to this main task - and under the policy direction of SACLANT - the Centre also renders scientific and technical assistance to the individual NATO nations.

This document is approved for public release.
Distribution is unlimited

SACLANT Undersea Research Centre
Viale San Bartolomeo 400
19138 San Bartolomeo (SP), Italy

tel: +39 0187 527 (1) or extension
fax: +39 0187 527 700

e-mail: library@saclantc.nato.int

NORTH ATLANTIC TREATY ORGANIZATION